Variational method for finding periodic orbits in a general flow

Yueheng Lan* and Predrag Cvitanović†
Center for Nonlinear Science, School of Physics,
Georgia Institute of Technology, Atlanta 30332-0430, U.S.A
(Dated: February 8, 2008)

A variational principle for determining unstable periodic orbits of flows as well as unstable spatio-temporally periodic solutions of extended systems is proposed and implemented. An initial loop approximating a periodic solution is evolved in the space of loops toward a true periodic solution by a minimization of local errors along the loop. The "Newton descent" partial differential equation that governs this evolution is an infinitesimal step version of the damped Newton-Raphson iteration. The feasibility of the method is demonstrated by its application to the Hénon-Heiles system, the circular restricted three-body problem, and the Kuramoto-Sivashinsky system in a weakly turbulent regime.

PACS numbers: 95.10.Fh, 02.70.Bf, 47.52.+j, 05.45.+a
Keywords: periodic orbits, variational methods, spatio-temporal chaos, turbulence, cost function minimization, Kuramoto-Sivashinsky system, restricted three-body problem, Hénon-Heiles system

I. INTRODUCTION

The periodic orbit theory of classical and quantum chaos [1, 2] is one of the major advances in the study of long-time behavior of chaotic dynamical systems. The theory expresses all long time averages over chaotic dynamics in terms of cycle expansions [3], sums over periodic orbits (cycles) ordered hierarchically according to the orbit length, stability, or action. If the symbolic dynamics is known, and the flow is hyperbolic, longer cycles are shadowed by the shorter ones, and cycle expansions converge exponentially or even super-exponentially with the cycle length [4].

A variety of methods for determining all periodic orbits up to a given length have been devised and successfully implemented for low-dimensional systems [5, 6, 7, 8, 9]. For more complex dynamics, such as turbulent flows [10], nonlinear waves [11], or quantum fields [12, 13] with high (or infinite) dimensional phase spaces and complicated dynamical behavior, many of the existing methods become unfeasible in practice. In the most computationally demanding calculation carried out so far, Kawahara and Kida [14] have found two periodic solutions in a $15,422$-dimensional discretization of a turbulent plane Couette flow. The topology of high-dimensional flows is hard to visualize, and even with a decent starting guess for the shape and location of a periodic orbit, methods like the Newton-Raphson method are likely to fail. In ref. [15] we have argued that variational, cost-function minimization methods offer a robust alternative. Here we derive, implement and discuss in detail one such new variational method for finding periodic orbits in general flows, and specifically high-dimensional flows.

In essence, any numerical algorithm for finding periodic orbits is based on devising a new dynamical system which possesses the desired orbit as an attracting fixed point with a sizable basin of attraction. Beyond that, there is much freedom in constructing such system.

For example, the multipoint shooting method eliminates the long-time exponential instability of unstable orbits by splitting an orbit into a number of short segments, each with a controllable expansion rate. The multiple shooting combined with the Newton-Raphson method is an efficient tool for locating periodic orbits of maps [16]. A search for periodic orbits of a continuous time flow can be reduced to a multiple shooting search for periodic orbits of a set of maps by constructing a set of phase space Poincaré sections such that an orbit leaving one section reaches the next one in a qualitatively predictable manner, without traversing other sections along the way. In turbulent, high-dimensional flows such sequences of sections are hard to come by. One solution might be a large set of Poincaré sections, with the intervening flight segments short and controllable.

Here we follow a different strategy, and discard Poincaré sections altogether; we replace maps between spatially fixed Poincaré sections, by maps induced by discretizing the time evolution into small time steps. For sufficiently small time steps such maps are small deformations of identity. We distribute many points along a smooth loop L, our initial guess of a cycle location and its topological layout. If both the time steps and the loop deformations are taken infinitesimal, a partial differential equation governs the "Newton descent", a fictitious time flow of a trial loop L into a genuine cycle p, with exponential convergence in the fictitious time variable. We then use methods developed for solving PDEs to get the solution. Stated succinctly, the idea of our method is to make an informed rough guess of what the desired cycle looks like globally, and then use a variational method to drive the initial guess towards the exact solution. For robustness, we replace the guess of a single orbit point by a guess of an entire orbit. For numerical safety we replace

*Electronic address: gte158y@prism.gatech.edu
†Electronic address: predrag.cvitanovic@physics.gatech.edu

the Newton-Raphson iteration by the "Newton descent", a differential flow that minimizes a cost function computed as deviation of the approximate flow from the true flow along a smooth loop approximation to a cycle.

In sect. II we derive the partial differential equation which governs the evolution of an initial guess loop toward a cycle and the corresponding cost function. An extension of the method to Hamiltonian systems and systems with higher time derivatives is presented in sect. III. Simplifications due to symmetries and details of our numerical implementation of the method are discussed in sect. IV. In sect. V we test the method on the Hénon-Heiles system, the restricted three body problem, and a weakly turbulent Kuramoto-Sivashinsky system. We summarize our results and discuss possible improvements of the method in sect. VI.

II. THE NEWTON DESCENT METHOD IN LOOP SPACE

A. A variational equation for the loop evolution

A periodic orbit is a solution (x, T), $x \in \mathbb{R}^d$, $T \in \mathbb{R}$ of the *periodic orbit condition*

$$f^T(x) = x\,, \qquad T > 0 \tag{1}$$

for a given flow or discrete time mapping $x \mapsto f^t(x)$. Our goal is to determine periodic orbits of flows defined by first order ODEs

$$\frac{dx}{dt} = v(x)\,, \qquad x \in \mathcal{M} \subset \mathbb{R}^d\,, \qquad (x, v) \in \mathbf{T}\mathcal{M} \tag{2}$$

in many (even infinitely many) dimensions d. Here $\mathcal{M}$ is the phase space (or state space) in which evolution takes place, $\mathbf{T}\mathcal{M}$ is the tangent bundle [17], and the vector field $v(x)$ is assumed to be smooth (sufficiently differentiable) almost everywhere.

We make our initial guess at the shape and the location of a cycle p by drawing a loop L, a smooth, differentiable closed curve $\tilde{x}(s) \in L \subset \mathcal{M}$, where s is a loop parameter. As the loop is periodic, we find it convenient to restrict s to $[0, 2\pi]$, with the periodic condition $\tilde{x}(s) = \tilde{x}(s + 2\pi)$. Assume that L is close to the true cycle p, pick N pairs of nearby points along the loop and along the cycle

$$\begin{aligned} \tilde{x}_n &= \tilde{x}(s_n)\,, \qquad 0 \le s_1 < \ldots < s_N < 2\pi\,, \\ x_n &= x(t_n)\,, \qquad 0 \le t_1 < \ldots < t_N < T_p\,, \end{aligned} \tag{3}$$

and denote by $\delta\tilde{x}_n$ the deviation of a point x_n on the periodic orbit p from the nearby point $\tilde{x}_n$,

$$x_n = \tilde{x}_n + \delta\tilde{x}_n\,.$$

The deviations $\delta\tilde{x}$ are assumed small, vanishing as L approaches p.

The orientation of the s-velocity vector tangent to the loop L

$$\tilde{v}(\tilde{x}) = \frac{d\tilde{x}}{ds}$$

is intrinsic to the loop, but its magnitude depends on the (still to be specified) parametrization s of the loop.

At each loop point $\tilde{x}_n \in L$ we thus have two vectors, the loop tangent $\tilde{v}_n = \tilde{v}(\tilde{x}_n)$ and the flow velocity $v_n = v(\tilde{x}_n)$. Our goal is to deform L until the directions of $\tilde{v}_n$ and v_n coincide for all $n = 1, \ldots, N$, $N \to \infty$, that is $L = p$. To match their magnitude, we introduce a local time scaling factor

$$\lambda(s_n) \equiv \Delta t_n / \Delta s_n\,, \tag{4}$$

where $\Delta s_n = s_{n+1} - s_n, n = 1, \ldots, N-1\,, \Delta s_N = 2\pi - (s_N - s_1)$, and likewise for Δt_n. The scaling factor $\lambda(s_n)$ ensures that the loop increment Δs_n is proportional to its counterpart $\Delta t_n + \delta t_n$ on the cycle when the loop L is close to the cycle p, with $\delta t_n \to 0$ as $L \to p$.

Let $x(t) = f^t(x)$ be the state of the system at time t obtained by integrating (2), and $J(x, t) = dx(t)/dx(0)$ be the corresponding Jacobian matrix obtained by integrating

$$\frac{dJ}{dt} = AJ\,, \quad A_{ij} = \frac{\partial v_i}{\partial x_j}\,, \qquad \text{with } J(x, 0) = \mathbb{1}\,. \tag{5}$$

Since the point $x_n = \tilde{x}_n + \delta\tilde{x}_n$ is on the cycle,

$$f^{\Delta t_n + \delta t_n}(\tilde{x}_n + \delta\tilde{x}_n) = \tilde{x}_{n+1} + \delta\tilde{x}_{n+1}\,. \tag{6}$$

Linearization

$$f^{\delta t}(x) \approx x + v(x)\delta t\,, \qquad f^t(x + \delta x) \approx x(t) + J(x, t)\delta x\,,$$

of (6) about the loop point $\tilde{x}_n$ and the time interval Δt_n to the next cycle point leads to the multiple shooting Newton-Raphson equation, for any step size Δt_n:

$$\delta\tilde{x}_{n+1} - J(\tilde{x}_n, \Delta t_n)\delta\tilde{x}_n - v_{n+1}\delta t_n = f^{\Delta t_n}(\tilde{x}_n) - \tilde{x}_{n+1}\,. \tag{7}$$

Provided that the initial guess is sufficiently good, the Newton-Raphson iteration of (7) generates a sequence of loops L with a decreasing cost function [15]

$$F^2(\tilde{x}) \equiv \frac{N}{(2\pi)^2} \sum_{i=1}^{N} (f^{\Delta t_n}(\tilde{x}_n) - \tilde{x}_{n+1})^2, \qquad \tilde{x}_{N+1} = \tilde{x}_1\,. \tag{8}$$

The prefactor $N/(2\pi)^2$ makes the definition of F^2 consistent with (13) in the $N \to \infty$ limit. If the flow is locally strongly unstable, the neighborhood in which the linearization is valid could be so small that the full Newton step would overshoot, rendering F^2 bigger rather than smaller. In this case the step-reduced, damped Newton method is needed. As proved in ref. [18], under conditions satisfied here, F^2 decreases monotonically if appropriate

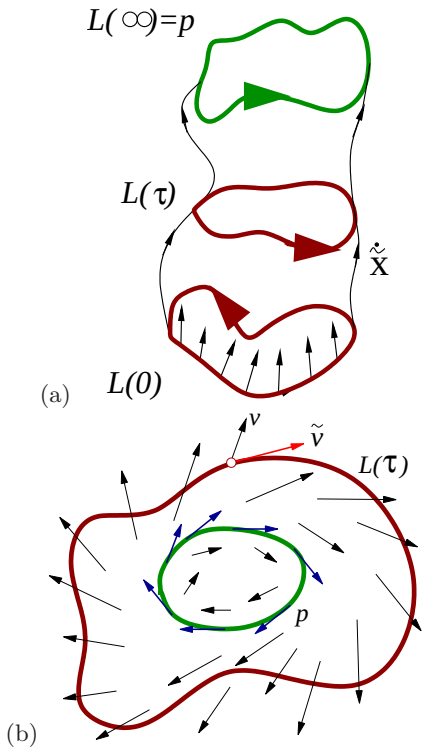

FIG. 1: (a) An annulus $L(\tau)$ swept by the Newton descent flow $d\tilde{x}/d\tau$, connecting smoothly the initial loop $L(0)$ to the periodic orbit $p = L(\infty)$. (b) In general the loop velocity field $\tilde{v}(\tilde{x})$ does not coincide with $\lambda v(\tilde{x})$; for a periodic orbit p, it does so at every $x \in p$.

step size is taken. If infinitesimal steps are taken, decrease of F^2 is ensured. We parametrize such continuous deformations of the loop by a *fictitious time* τ.

We fix Δs_n and proceed by $\delta\tau$ each step of the iteration, that is, multiply the right hand side of (7) by $\delta\tau$. According to (4), the change of Δt_n with respect to τ is equal to $\delta t_n = \frac{\partial\lambda}{\partial\tau}(s_n,\tau)\delta\tau\Delta s_n$. As $\delta\tilde{x}_n = \frac{\partial}{\partial\tau}\tilde{x}(s_n,\tau)\delta\tau$, dividing both sides of (7) by $\delta\tau$ yields

$$\frac{d\tilde{x}_{n+1}}{d\tau} - J(\tilde{x}_n,\Delta t_n)\frac{d\tilde{x}_n}{d\tau} - v_{n+1}\frac{\partial\lambda}{\partial\tau}(s_n,\tau)\Delta s_n = f^{\Delta t_n}(\tilde{x}_n) - \tilde{x}_{n+1}\,. \qquad (9)$$

In the $N \to \infty$ limit, the stepsizes $\Delta s_n, \Delta t_n = O(1/N) \to 0$, and we have

$$v_{n+1} \approx v_n\,, \qquad \tilde{x}_{n+1} \approx \tilde{x}_n + \tilde{v}_n\Delta s_n\,,$$
$$J(\tilde{x}_n,\Delta t_n) \approx 1 + A(\tilde{x}_n)\Delta t_n\,, \qquad f^{\Delta t_n}(\tilde{x}_n) \approx \tilde{x}_n + v_n\Delta t_n\,.$$

Substituting into (9) and using the scaling relation (4), we obtain

$$\frac{\partial^2\tilde{x}}{\partial s\partial\tau} - \lambda A\frac{\partial\tilde{x}}{\partial\tau} - v\frac{\partial\lambda}{\partial\tau} = \lambda v - \tilde{v}. \qquad (10)$$

This PDE, which describes the evolution of a loop $L(\tau)$ toward a periodic orbit p, is the central result of this paper. The family of loops so generated is parametrized by $\tilde{x} = \tilde{x}(s,\tau) \in L(\tau)$, where s denotes the position along the loop, and the fictitious time τ parametrizes the deformation of the loop, see Fig. 1(a). We refer to this infinitesimal step version of the damped Newton-Raphson method as the "Newton descent".

The important feature of this equation is that a decreasing cost functional exists. Rewriting (10) as

$$\frac{\partial}{\partial\tau}(\tilde{v} - \lambda v) = -(\tilde{v} - \lambda v)\,, \qquad (11)$$

we have

$$\tilde{v} - \lambda v = e^{-\tau}(\tilde{v} - \lambda v)|_{\tau=0}\,, \qquad (12)$$

so the fictitious time τ flow decreases the cost functional

$$F^2[\tilde{x}] = \frac{1}{2\pi}\oint_{L(\tau)} ds\ (\tilde{v}(\tilde{x}) - \lambda v(\tilde{x}))^2 \qquad (13)$$

monotonically as the loop evolves toward the cycle.

At each iteration step the differences of the loop tangent velocities and the dynamical flow velocities are reduced by the Newton descent. As $\tau \to \infty$, the fictitious time flow alignes the loop tangent $\tilde{v}$ with the dynamical flow vector $\tilde{v} = \lambda v$, and the loop $\tilde{x}(s,\tau) \in L(\tau)$, see Fig. 1(b), converges to a genuine periodic orbit $p = L(\infty)$ of the dynamical flow $\dot{x} = v(x)$. Once the cycle p is reached, by (4), $\lambda(s,\infty) = \frac{dt}{ds}(\tilde{x}(s,\infty))$, and the cycle period is given by

$$T_p = \int_0^{2\pi} \lambda(\tilde{x}(s,\infty))ds\,.$$

Of course, as at this stage we have already identified the cycle, we may pick instead an initial point on p and calculate the period by a direct integration of the dynamical equations (2).

B. Marginal directions and accumulation of loop points

Numerically, two perils lurk in a direct implementation of the Newton descent (10).

First, when a cycle is reached, it remains a cycle under a cyclic permutation of the representative points, so on the cycle the operator

$$\bar{A} = \frac{\partial}{\partial s} - \lambda A$$

has a marginal eigenvector $v(\tilde{x}(s))$ with eigenvalue 0. If λ is fixed, as the loop approaches the cycle, (10) approaches its limit

$$\bar{A}\,\frac{\partial x}{\partial\tau} = 0\,.$$

Therefore, on the cycle, the operator $\bar{A}^{-1}$ becomes singular and the numerical woes arise.

The second potential peril hides in the freedom of choosing the loop (re-)parametrization. Since s is related to the time t by the yet unspecified factor $\lambda(s,\tau)$, uneven distributions of the sampling points over the loop L could arise, with the numerical discretization points $\tilde{x}_n$ clumping densely along some segments of L and leaving big gaps elsewhere, thus degrading the numerical smoothness of the loop.

We remedy these difficulties by imposing constraints on (10). In our calculation for Kuramoto-Sivashinsky system of sect. V, the first difficulty is dealt with by introducing one Poincaré section, for example, by fixing one coordinate of one of the sampling points, $\tilde{x}_1(s_1,\tau) = const$. This breaks the translational invariance along the cycle. Other types of constraints might be better suited to a specific problem at hand. For example, we can demand that the average displacement of the sampling points along the loop vanishes, thus avoiding a spiraling descent towards the desired cycle.

We deal with the second potential difficulty by choosing a particularly simple loop parametrization. So far, the parametrization s is arbitrary and there is much freedom in choosing the best one for our purposes. We pick $s-$ and $\tau-$independent constant scaling $\lambda(s,\tau) = \lambda$. With uniform grid size $\Delta s_n = \Delta s$ and fixed λ, the loop parameter $s = t/\lambda$ is proportional to time t, and the discretization (10) distributes the sampling points along the loop evenly in time. As the loop approaches a cycle, $\frac{\partial \tilde{x}}{\partial \tau}$ is numerically obtainable from (10), and on the cycle the period is given by $T_p = 2\pi\lambda$.

Even though this paper focuses on searches for periodic orbits, the Newton descent is a general method. With appropriate modifications of boundary conditions and scaling of time, (10) can be adapted to determination of homoclinic or heteroclinic orbits between equilibrium points or periodic orbits of a flow, or more general boundary value problems. Applied to 2-point boundary value problems, Newton descent is similar to the quasilinearization [19] but has the advantage that the free parameter $\lambda(s,\tau)$ is available for adjusting scales in the problem and that searches can be restricted to phase space submanifolds of interest. A simple example of a restriction to a submanifold are searches for cycles of a given energy, constrained to the $H(q,p) = E$ energy shell in the phase space of a Hamiltonian system. Furthermore, as we shall show now, the symplectic structure of Hamilton's equations greatly reduces the dimensionality of the submanifold that we need to consider.

III. EXTENSIONS OF NEWTON DESCENT

In classical mechanics particle trajectories are also solutions of a variational principle, the Hamilton's variational principle. For example, one can determine a periodic orbit of a billiard by wrapping around a rubber band of a roughly correct topology, and then moving the points along the billiard walls until the length (that is, the action) of the rubber band is extremal (maximal or minimal under infinitesimal changes of the boundary points). Note that the extremization of action requires only D configuration coordinate variations, not the full $2D$-dimensional phase space variations.

Can we exploit this property of the Newtonian mechanics to reduce the dimenionality of our variational calculations? The answer is yes, and easiest to understand in terms of the Hamilton's variational principle which states that classical trajectories are extrema of the Hamilton's principal function (or, for fixed energy E, the action $S = R + Et$)

$$R(q_1,t_1;q_0,t_0) = \int_{t_0}^{t_1} dt\, \mathcal{L}(q(t),\dot{q}(t),t)\,,$$

where $\mathcal{L}(q,\dot{q},t)$ is the Lagrangian. Given a loop $L(\tau)$ we can compute not only the tangent "velocity" vector $\tilde{v}$, but also the local loop curvature or "acceleration" vector

$$\tilde{a} = \frac{\partial^2 \tilde{x}}{\partial s^2}\,,$$

and indeed, as many s derivatives as needed. Matching the dynamical acceleration $a(\tilde{x})$ (assumed to be functions of $\tilde{x}$ and $v(\tilde{x})$) with the loop "acceleration" $\tilde{a}(\tilde{x})$ results in a new cost function and the corresponding PDE (11) for the evolution of the loop

$$\frac{\partial}{\partial \tau}(\tilde{a} - \lambda^2 a) = -(\tilde{a} - \lambda^2 a)\,.$$

We use λ^2 instead of λ in order to keep the notation consistent with (4), that is $t = \lambda\, s$. Expressed in terms of the loop variables $\tilde{x}(s)$, the above equation becomes

$$\frac{\partial^3 \tilde{x}}{\partial^2 s \partial \tau} - \lambda \frac{\partial a}{\partial v}\frac{\partial^2 \tilde{x}}{\partial s \partial \tau} - \lambda^2 \frac{\partial a}{\partial \tilde{x}}\frac{\partial \tilde{x}}{\partial \tau} + \left(\frac{\partial a}{\partial v}\frac{\partial \tilde{x}}{\partial s} - 2\lambda a\right)\frac{\partial \lambda}{\partial \tau} = \lambda^2 a - \tilde{a}\,, \tag{14}$$

where $v = \frac{\partial \tilde{x}}{\lambda \partial s}$. Although (14) looks more complicated than (10), in numerical fictitious time integrations, we are rewarded by having to keep only half of the phase space variables.

More generally, if a differential equation has the form:

$$x^{(m)} = f(x, x^{(1)}, \cdots, x^{(m-1)})\,, \tag{15}$$

where $x^{(k)} = \frac{d^k x}{dt^k}$, $k = 1, \cdots, m$ and $x \in \mathbb{R}^d$, the same technique can be used to match the highest derivatives $\lambda^m x^{(m)}$ and $\tilde{x}^{(m)}$,

$$\frac{\partial}{\partial \tau}(\tilde{x}^{(m)} - \lambda^m x^{(m)}) = -(\tilde{x}^{(m)} - \lambda^m x^{(m)})\,,$$

with $\tilde{x}^{(m)} = \frac{\partial^m}{\partial s^m}\tilde{x}(s)$ calculated directly from $\tilde{x}(s)$ on the loop by differentiation. In loop variables $\tilde{x}(s)$ we have,

$$\frac{\partial^{m+1} \tilde{x}}{\partial s^m \partial \tau} - \lambda^m \sum_{k=0}^{m} \frac{\partial f}{\partial x^{(k)}} \cdot \frac{\partial}{\partial \tau}\frac{\partial^k \tilde{x}}{\lambda^k \partial s^k} - m\lambda^{m-1}\tilde{x}^{(m)}\frac{\partial \lambda}{\partial \tau} = \lambda^m x^{(m)} - \tilde{x}^{(m)}\,, \tag{16}$$

where $x = x^{(0)}$ and $\tilde{x}^{(k)} = \frac{\partial^k \tilde{x}}{\lambda^k \partial^k s}$, $k = 1, \cdots, m-1$ are assumed. Conventionally, (15) is converted to a system of md first order differential equations, whose discretized derivative (see (17) below) are banded matrices with band width of $5md$. Using (16), we only need d equations for the same accuracy and the corresponding band width is $(m+4)d$. The computing load has been greatly reduced, the more so the larger m is. Nevertheless, choice of a good initial loop guess and visualization of the dynamics are always aided by a plot of the orbit in the full md-dimensional phase space, where loops cannot self-intersect and topological features of the flow is exhibited more clearly.

IV. IMPLEMENTATION OF NEWTON DESCENT

As the loop points satisfy a periodic boundary condition, it is natural to employ truncated discrete Fast Fourier Transforms (FFT) in numerical integrations of (10). Since we are interested only in the final, stationary cycle p, the accuracy of the fictitious time integration is not crucial; all we have to ensure is the smoothness of the loop throughout the integration. The Euler integration with fairly large time steps $\delta\tau$ suffices. The computationally most onerous step in implementation of the Newton descent is the inversion of large matrix $\bar{A}$ in (10). When the dimension of the dynamical phase space of (2) is high, the inversion of $\bar{A}$ needed to get $\frac{\partial \tilde{x}}{\partial \tau}$ takes most of the integration time, making the evolution extremely slow. This problem is partially solved if the finite difference methods are used. The large matrix $\bar{A}$ then becomes sparse and the inversion can be done far more quickly.

A. Numerical implementation

In a discretization of a loop, numerical stability requires accurate discretization of loop derivatives such as

$$\tilde{v}_n \equiv \left.\frac{\partial \tilde{x}}{\partial s}\right|_{\tilde{x}=\tilde{x}(s_n)} \approx (\hat{D}\tilde{x})_n \,.$$

In our numerical work we use the four-point approximation [20],

$$\hat{D} = \frac{1}{12h}\begin{pmatrix} 0 & 8 & -1 & & & & & & 1 & -8 \\ -8 & 0 & 8 & -1 & & & & & & 1 \\ 1 & -8 & 0 & 8 & -1 & & & & & \\ & & & & & \cdots & & & & \\ & & & & & & 1 & -8 & 0 & 8 & -1 \\ -1 & & & & & & & 1 & -8 & 0 & 8 \\ 8 & -1 & & & & & & & 1 & -8 & 0 \end{pmatrix} \tag{17}$$

where $h = 2\pi/N$. Here, each entry represents a $[d \times d]$ matrix, $8 \to 8\mathbb{1}$, *etc.*, with blank spaces are filled with zeros. The two $[2d \times 2d]$ matrices

$$M_1 = \begin{pmatrix} \mathbb{1} & -8\mathbb{1} \\ 0 & \mathbb{1} \end{pmatrix}, \qquad M_2 = \begin{pmatrix} -\mathbb{1} & 0 \\ 8\mathbb{1} & -\mathbb{1} \end{pmatrix},$$

located at the top-right and bottom-left corners take care of the periodic boundary condition.

The discretized version of (10) with a fictitious time Euler step $\delta\tau$ is

$$\begin{pmatrix} \hat{A} & \hat{v} \\ \hat{a} & 0 \end{pmatrix} \begin{pmatrix} \delta\hat{x} \\ \delta\lambda \end{pmatrix} = \delta\tau \begin{pmatrix} \lambda\hat{v} - \hat{\tilde{v}} \\ 0 \end{pmatrix}, \tag{18}$$

where

$$\hat{A} \;=\; \hat{D} + \mathrm{diag}[A_1, A_2, \cdots, A_N]\,,$$

with $A_n = A(\tilde{x}(s_n))$ defined in (5), and

$$\begin{aligned} \hat{v} \;&=\; (v_1, v_2, \cdots, v_N)\,, \qquad && \text{with } v_n = v(\tilde{x}(s_n))\,, \\ \hat{\tilde{v}} \;&=\; (\tilde{v}_1, \tilde{v}_2, \cdots, \tilde{v}_N)\,, \qquad && \text{with } \tilde{v}_n = \tilde{v}(\tilde{x}(s_n))\,, \end{aligned}$$

are the two vector fields that we want to match everywhere along the loop. $\hat{a}$ is an Nd dimensional row vector which imposes the constraint on the coordinate variations $\delta\hat{x} = (\delta\tilde{x}_1, \delta\tilde{x}_2, \cdots, \delta\tilde{x}_N)$. The discretized Newton descent (18) is an infinitesimal time step variant of the multipoint (Poincaré section) shooting equation for flows [16]. In order to solve for the deformation of the loop coordinates and period, $\delta\hat{x}$ and $\delta\lambda$, we need to invert the $[(N\,d+1)\times(N\,d+1)]$ matrix on the left hand side of (18).

In our numerical work, this matrix is inverted using the banded LU decomposition on the embedded band-diagonal matrix, and the Woodbury formula [21] on the cyclic, border terms. The LU decomposition takes most of the computational time and considerably slows down the fictitious time integration. We speed up the integration by a new inversion scheme which relies on the smoothness of the flow in the loop space. It goes as follows. Once we have the LU decomposition at one step, we use it to approximately invert the matrix in the next step, with accurate inversion achieved by the iterative approximate inversions [21]. In our applications we find that a single LU decomposition can be used for many $\delta\tau$ evolution steps. The further we go, the more iterations at each step are needed to implement the inversion. After the number of such iterations exceeds some given fixed maximum number, we perform another LU decomposition and proceed as before. The number of integration steps following one decomposition is an indication of the smoothness of the evolution, and we adjust accordingly the integration step size $\delta\tau$: the greater the number, the bigger the step size. As the loop approaches a cycle, the evolution becomes so smooth that the step size can be brought all the way up to $\delta\tau = 1$, the full undamped Newton-Raphson iteration step. In practice, one can start with a small but reasonable number of points, in order to get a coarse solution of relatively low accuracy. After achieving that, the refined guess loop can be constructed by interpolating more points, and proceed with for a more accurate calculation in which $\delta\tau$ can be set as large as the full Newton step $\delta\tau = 1$, recovering the rapid quadratic convergence of the Newton-Raphson method.

It is essential that the smoothness of the loop is maintained throughout the calculation. We monitor the smoothness by checking the Fourier spectrum of $\tilde{x}(\cdot,\tau)$. An unstable difference scheme for loop derivatives might lead to unbounded sawtooth oscillations [22]. A heuristic local linear stability analysis (described in [23]) indicates that our scheme is stable, and that the high frequency components do not generate instabilities.

B. Initialization and convergence

As in any other method, a qualitative understanding of the dynamics is a prerequisite to successful cycle searches. We start by numerical integration with the dynamical system (2). Numerical experiments reveal regions where a trajectory spends most of its life, giving us the first hunch as to how to initialize a loop. We take the FFT of some nearly recurred orbit segment and keep only the lowest frequency components. The inverse Fourier transform back to the phase space yields a smooth loop that we use as our initial guess. Since any generic orbit segment is not closed and might exhibit large gaps, the Gibbs phenomenon can take the initial loop so constructed quite far away from the region of interest. We deal with this problem by manually deforming the orbit segment into a closed loop before performing the FFT. Searching for longer cycles with multiple circuits requires more delicate initial conditions. The hope is that a few short cycles can help us establish an approximate symbolic dynamics, and guess for longer cycles can be constructed by cutting and glueing the short, known ones. For low dimensional systems, such methods yield quite good systematic initial guesses for longer cycles [24].

An alternative way to initialize the search is by utilizing adiabatic deformations of dynamics, or the homotopy evolution [25]. If the dynamical system (2) depends on a parameter μ, short cycles might survive as μ varies passing through a family of dynamical systems, giving in the process birth to new cycles through sequences of bifurcations. Most short unstable cycles vary little for small changes of μ. So, a cycle existing for parameter value μ_1 can be chosen as the initial trial loop for a nearby cycle surviving a small change $\mu_1 \to \mu_2$. In practice, one or two iterations often suffice to find the new cycle.

A good choice of the initial loop significantly expedites the computation, but there are more reasons why good initial loops are crucial. First of all, if we break the translational invariance by imposing a constraint such as $\tilde{x}_1(s_1,\tau) = c$, we have to make sure that both the initial loop and the desired cycle intersects this Poincaré plane. Hence, the initial loop cannot be wildly different from the desired cycle. Second, in view of (12), the loop always evolves towards a local minimum of the cost functional (13), with discretization points moving along the $\tilde{v}-\lambda v$ fixed direction, determined by the initial condition. If the local minimum corresponds to a zero of the cost functional, we obtain a true cycle of the dynamical flow (2). However, if the value of cost functional is not equal to zero at the minimum while the gradient is zero, (18) yields a singular matrix $\hat{A}$. In such cases the search has to be abandoned and restarted with a new initial loop guess. In the periodic orbit searches of sect. V starting with blind initial guesses (guesses unaided by a symbolic dynamics partition), such local minima were encountered in about 30% of cases.

C. Symmetry considerations

The system under consideration often possesses certain symmetries. If this is the case, the symmetry should be both be feared for possible marginal eigen-directions, and be embraced as a guide to possible simplifications of the numerical calculation.

If the dynamical system equations (2) are invariant under a discrete symmetry, the concept of fundamental domain [5, 26] can be utilized to reduce the length of the initial loop when searching for a cycle of a given symmetry. In this case, we need discretize only an irreducible segment of the loop, decreasing significantly the dimensionality of the loop representation. Other parts of the loop are replicated by symmetry operations, with the full loop tiled by copies of the fundamental domain segment. The boundary conditions are not periodic any longer, but all that we need to do is modify the cyclic terms. Instead of using M_1 and M_2 in (17), we use $M_1 Q$ and $M_2 Q^{-1}$, where Q is the relevant symmetry operation that maps the fundamental segment to the neighbor that precedes it. In this way, a fraction of the points represent the cycle with the same accuracy, speeding up the search considerably.

If a continuous symmetry is present, it may complicate the situation at first glance but becomes something that we can take advantage of after careful checking. For example, for a Hamiltonian system unstable cycles may form continuous families [27, 28], with one or more members of a family belonging to a given constant energy surface. In order to cope with the marginal eigendirection associated with such continous family, we search for a cycle on a particular energy surface by replacing the last row of equation (18) by an energy shell constraint [16]. We put one point of the loop, say $\tilde{x}_2$, on the constant energy surface $H(\tilde{x}) = E$, and impose the constraint $\bigtriangledown H(\tilde{x}_2)\cdot\delta\tilde{x}_2 = 0$, so as to keep $\tilde{x}_2$ on the surface for all τ. The integration of (10) then automatically brings all other loop points to the same energy surface. Alternatively, we can look for a cycle of given fixed period T by fixing λ and dropping the constraint in the bottom line of (18). These two approaches are conjugate to each other, both needed in applications. In most cases, they are equivalent. One exception is the harmonic oscillator for which the oscillations have identical period but different energy. Note that in both cases the translational invariance is restored, as we have discarded the Poincaré section condition of sect. II B. As explained in [6], this

causes no trouble in numerical calculations.

V. APPLICATIONS

We have checked that the iteration of (18) yields quickly and robustly the short unstable cycles for standard models of low-dimensional dissipative flows such as the Rössler system [29]. A more daunting challenge are searches for cycles in Hamiltonian flows, and searches for spatio-temporally periodic solutions of PDEs. In all numerical examples that follow, the convergence condition is $F^2 < 10^{-5}$.

A. Hénon-Heiles system and restricted three-body problem

First, we test the Hamiltonian version of the Newton descent derived in Sect. III by applying the method to two Hamiltonian systems, both with two degrees of freedom. In both cases, our initial loop guesses are rather arbitrary combinations of trigonometric functions. Nevertheless, the observed convergence is fast.

The Hénon-Heiles system [30] is a standard model in celestial mechanics, described by the Hamiltonian

$$H = \frac{1}{2}(p_x^2 + p_y^2 + x^2 + y^2) + x^2 y - \frac{y^3}{3}\,. \qquad (19)$$

It has a time reversal symmetry and a three-fold discrete spatial symmetry. Figure 2 shows a typical application of (14), with the Newton descent search restricted to the configuration space. The initial loop, Fig. 2(a), is a rather coarse initial guess. We fix arbitrarily the scaling $\lambda = 2.1$, that is, we search for a cycle p of the fixed period $T_p = 13.1947$, with no constraint on the energy. Figure 2(b) shows the cycle found by the Newton descent, with energy $E = 0.1794$, and the full discrete symmetry of the Hamiltonian. This cycle persists adiabatically for a small range of values of λ; with λ changed much, the Newton descent takes the same initial loop into other cycles. Figure 2(c) verifies that the cost functional F^2 decreases exponentially with slope -2 throughout the $\tau = [0, 10]$ integration interval, as predicted by (12). The points get more and more sparse as τ increases, because our numerical implementation adaptively chooses bigger and bigger step sizes $\delta\tau$.

In the Hénon-Heiles case, the accelerations a_x, a_y depend only on the configuration variables x, y. More generally, the accelerations could also depend on $\dot{x}, \dot{y}$. Consider as an example the equations of motion for the restricted three-body problem [31],

$$\begin{aligned} \ddot{x} &= 2\dot{y} + x - (1-\mu)\frac{x+\mu}{r_1^3} - \mu\frac{x-1+\mu}{r_2^3}\,, \\ \ddot{y} &= -2\dot{x} + y - (1-\mu)\frac{y}{r_1^3} - \mu\frac{y}{r_2^3}\,, \end{aligned} \qquad (20)$$

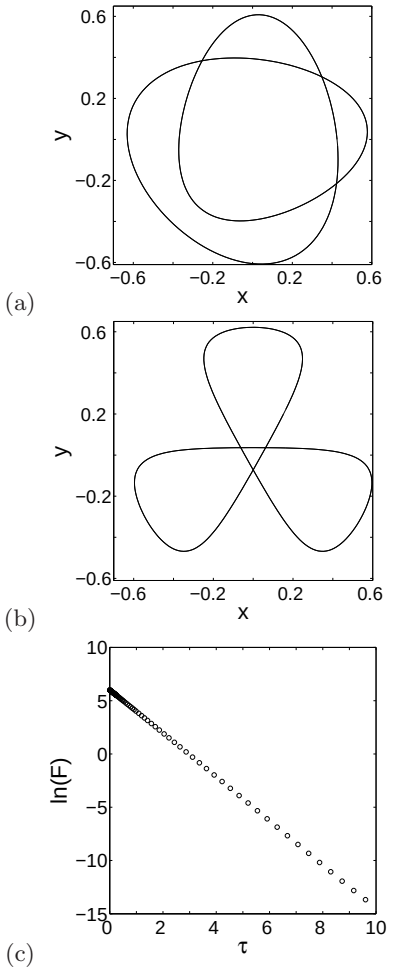

FIG. 2: The Hénon-Heiles system in a chaotic region: (a) An initial loop $L(0)$, and (b) the unstable periodic orbit p of period $T = 13.1947$ reached by the Newton descent (14). (c) The exponential decrease of the cost function, $\ln(F^2) \approx -2.0502\,\tau + 6.0214$.

where $r_1 = \sqrt{(x+\mu)^2 + y^2}$, $r_2 = \sqrt{(x-1+\mu)^2 + y^2}$. These equations describe the motion of a test particle in a rotating frame under the influence of the gravitational force of two heavy bodies with masses 1 and $\mu \ll 1$ fixed at $(-\mu, 0)$ and $(1-\mu, 0)$ in the (x, y) coordinate frame. The stationary solutions of (20) are called the Lagrange points, corresponding to a circular motion of the test particle in phase with the rotation of the heavy bodies. The periodic solutions in the rotating frame correspond to periodic or quasi-periodic motion of the test particle in the inertial frame. Figure 3 shows an initial loop and the cycle to which it converges, in the rotating frame. Although the cycle looks simple, the Newton descent requires advancing in small $\delta\tau$ steps in order for the initial loop to converge to it.

In order to successfully apply the Hamiltonian version of the Newton descent (14), we have to ensure that the test particle keeps a finite distance from the origin. If a cycle passes very close to one of the heavy bodies,

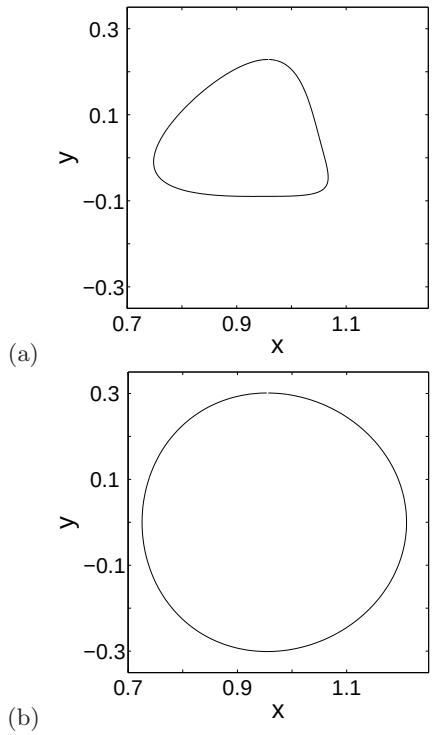

FIG. 3: (a) An initial loop $L(0)$, and (b) the unstable periodic orbit p of period $T_p = 2.7365$ reached by the Newton descent (14), the restricted three body problem (20) in the chaotic regime, $\mu = 0.04$.

the acceleration can become so large that our scheme of uniformly distributing the loop points in time might fail to represent the loop faithfully. Another distribution scheme is required in this case, for example, making the density of points proportional to the magnitude of acceleration.

B. Periodic orbits of Kuramoto-Sivashinsky system

The Kuramoto-Sivashinsky equation arises as an amplitude equation for interfacial instability in a variety of contexts [32, 33]. In 1-dimensional space, it reads

$$u_t = (u^2)_x - u_{xx} - \nu u_{xxxx}, \tag{21}$$

where ν is a "super-viscosity" parameter which controls the rate of dissipation and $(u^2)_x$ is the nonlinear convection term. As ν decreases, the system undergoes a series of bifurcations, leading to increasingly turbulent, spatio-temporally chaotic dynamics.

If we impose the periodic boundary condition $u(t, x + 2\pi) = u(t, x)$ and choose to study only the odd solutions $u(-x, t) = -u(x, t)$, $u(x, t)$ can be expanded in a discrete spatial Fourier series [24],

$$u(x,t) = i \sum_{k=-\infty}^{\infty} a_k(t) e^{ikx}, \tag{22}$$

where $a_{-k} = -a_k \in \mathbb{R}$. In terms of the Fourier components, PDE (21) becomes an infinite ladder of ODEs,

$$\dot{a_k} = (k^2 - \nu k^4) a_k - k \sum_{m=-\infty}^{\infty} a_m a_{k-m} \,. \tag{23}$$

In numerical simulations we work with the Galerkin truncations of the Fourier series since in the neighborhood of the strange attractor the magnitude of a_k decreases very fast with k, high frequency modes playing a negligible role in the asymptotic dynamics. In this way Galerkin truncations reduce the dynamics to a finite but large number of ODEs. We work with $d = 32$ dimensions in our numerical calculations. In ref. [24], multipoint shooting has been successfully applied to obtain periodic orbits close to the onset of spatiotemporal chaos ($\nu = 0.03$). In this regime, our method is so stable that big time steps $\delta\tau$ can be employed even at the initial guesses, leading to extremely fast convergence. We attribute this robustness to the simplicity of the structure of the attractor at high viscosity values.

The challenge comes with decreasing ν, with the dynamics turning more and more turbulent. Already at $\nu = 0.015$, the system is moderately turbulent and the phase space portraits of the dynamics reveal a complex labyrinth of "eddies" of different scales and orientations. While the highly unstable nature of orbits and intricate structure of the invariant set hinder applications of conventional cycle-search routines, in this setting our variational method shines through. We design rather arbitrary initial loops from numerical trajectory segments, and the calculation proceeds as before, except that now a small $\delta\tau$ has to be used initially to ensure numerical stability. Topologically different loops are very likely to result in different cycles, while some initial loop guesses my lead to local nonzero minima of the cost functional F^2. As explained in Sect. IV, in such cases the method diverges, and the search is restarted with a new inital loop guess.

Two initial loop guesses are displayed in Fig. 4, along with the two periodic orbits detected by the Newton descent. In discretization of the initial loops, each point has to be specified in all d dimensions; here the coordinates $\{a_1, a_2\}$ are picked so that topological similarity between initial and final loops is visually easy to identify. Other projections from $d = 32$ dimensions to subsets of 2 coordinates appear to make the identification harder, if not impossible. In both calculations, we molded segments of typical trajectories into smooth closed loops by the Fourier filtering method of Sect. IV. As the desired orbit becomes longer and more complex, more sampling points are needed to represent the loop. We use $N = 512$ points to represent the loop in the (a)-(b) case and $N = 1024$ points in the (c)-(d) case. The space-time evolution of $u(x, t)$ for these two unstable spatio-temporally periodic solutions is displayed in Fig. 5. As $u(x, t)$ is antisymmetric on $[-\pi, \pi]$, it suffices to display the solutions on the $x \in [0, \pi]$ interval.

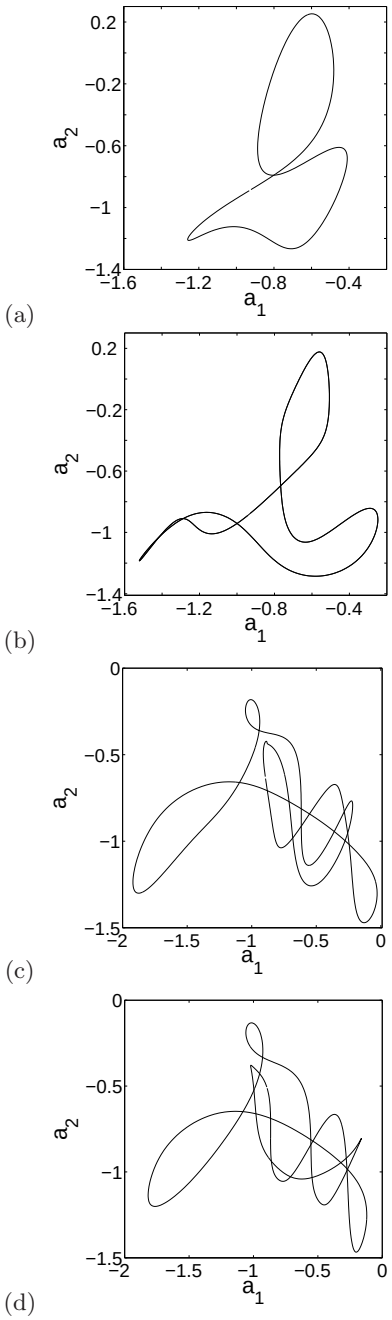

FIG. 4: The Kuramoto-Sivashinsky system in a spatiotemporally turbulent regime (viscosity parameter $\nu = 0.015$, $d = 32$ Fourier modes truncation). (a) An initial guess L_1, and (b) the periodic orbit p_1 of period $T_1 = 0.744892$ reached by the Newton descent. (c) Another initial guess L_2, and (d) the resulting periodic orbit p_2 of period $T_2 = 1.184668$.

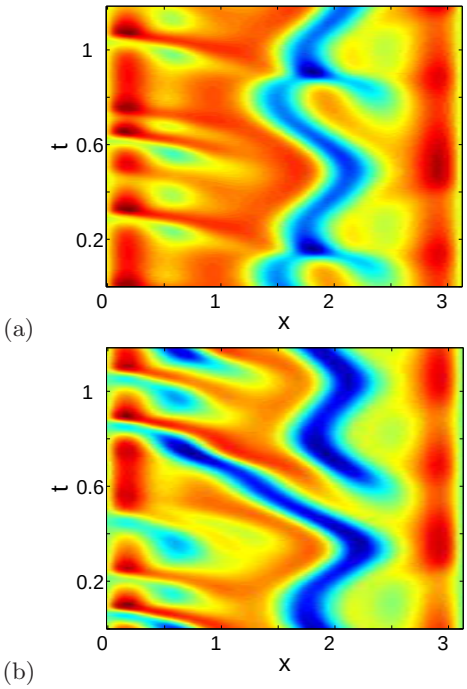

FIG. 5: Level plot of the space-time evolution $u(x,t)$ for the two spatiotemporally periodic solutions of Fig. 4: (a) the evolution of p_1, with start of a repeat after the cycle period $T_1 = 0.744892$, and (b) one full period $T_2 = 1.184668$ in the evolution of p_2.

VI. DISCUSSION

In order to cope with the difficulty of finding periodic orbits in high-dimensional chaotic flows, we have devised the *Newton descent method*, an infinitesimal step variant of the damped Newton-Raphson method. Our main result is the PDE (10) which solves the variational problem of minimizing the cost functional (13). This equation describes the fictitious time τ flow in the space of loops which decreases the cost functional at uniform exponential rate (see (12)). Variants of the method are presented for special classes of systems, such as Hamiltonian systems. An efficient integration scheme for the PDE is devised and tested on the Kuramoto-Sivashinsky system, the Hénon-Heiles system and the restricted three-body problem.

Our method uses information from a large number of points in phase space, with the global topology of the desired cycle protected by insistence on smoothness and a uniform discretization of the loop. The method is quite robust in practice.

The numerical results presented here are only a proof of principle. We do not know to what periodic orbit the flow (10) will evolve for a given dynamical system and a given initial loop. Empirically, the flow goes toward the "nearest" periodic orbit, with the largest topological resemblance. Each particular application will still require much work in order to elucidate and enumerate relevant topological structures. The hope is that the short spatiotemporally periodic solutions revealed by the Newton descent searches will serve as the basic building blocks for systematic investigations of chaotic and perhaps even "turbulent" dynamics.

Acknowledgments

We would like to thank Cristel Chandre for careful reading of the manuscript and numerous suggestions.

[1] D. Ruelle, *Statistical Mechanics, Thermodynamic Formalism* (Addison-Wesley, Reading MA, 1978).
[2] M. Gutzwiller, *Chaos in Classical and Quantum Mechanics* (Springer-Verlag, New York, 1990).
[3] R. Artuso, E. Aurell, and P. Cvitanović, Nonlinearity **3**, 325 (1990).
[4] H. Rugh, Nonlinearity **5**, 1237 (1992).
[5] P. Cvitanović, R. Artuso, R. Mainieri, G. Tanner, and G. Vattay, *Chaos: Classical and Quantum* (Niels Bohr Institute, Copenhagen, 2003), `www.nbi.dk/ChaosBook`.
[6] D. Viswanath, unpublished.
[7] F. Diakonos, P. Schmelcher, and O. Biham, Phys. Rev. Lett. **81**, 4319 (1998), `chao-dyn/9810022`.
[8] O. Biham and W. Wenzel, Phys. Rev. Lett. **63**, 819 (1989).
[9] R. L. Davidchack and Y.-C. Lai, Phys. Rev. E **60**, 6172 (1999).
[10] U. Frisch, *Turbulence* (Cambrige University Press, Cambrige, 1996).
[11] I. S. Aranson and L. Kramer, **74**, 99 (2002).
[12] T. S. Biró, S. G. Matinyan, and B. Müller, *Chaos and Gauge Field Theory* (World Scientific, Singapore, 1994).
[13] P. Cvitanović, Physica A **288**, 61 (2000), `nlin.CD/0001034`.
[14] G. Kawahara and S. Kida, J. Fluid Mech. **449**, 291 (2001).
[15] P. Cvitanović and Y. Lan, in *Proceed. of 10. Intern. Workshop on Multiparticle Production: Correlations and Fluctuations in QCD*, edited by N. Antoniou (World Scientific, Singapore, 2003), `nlin.CD/0308006`.
[16] F. Christiansen, chapter *"Fixed points, and how to get them"*, in ref. [5].
[17] V. I. Arnol'd, *Ordinary Differential Equations* (Springer, New York, 1992).
[18] H. B. Keller, *Numerical methods for two-point boundary-value problems* (Dover, New York, 1992).
[19] J. Stoer and R. Bulirsch, *Introduction to Numerical Analysis* (Springer-Verlag, New York, 1983).
[20] A. Brandenburg, in *Advances in nonlinear dynamos*, edited by A. Ferriz-Mas and M. Núñez (Taylor & Francis, London, 2003), `astro-ph/0109497`.
[21] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C* (Cambridge University Press, Cambridge, 1992).
[22] J. W. Thomas, *Numerical Partial Differential Equations* (Spring-Verlag, New York, 1995).
[23] B. Roždestvenskiĭ and N. Janenko, *Systems of Quasilinear Equations and Their Applications to Gas Dynamics* (AMS, Providence, 1983).
[24] F. Christiansen, P. Cvitanović, and V. Putkaradze, Nonlinearity **10**, 55 (1997).
[25] C. Conley, *Isolated Invariant Sets and the Morse Index* (AMS, Providence, 1978).
[26] P. Cvitanović and B. Eckhardt, Nonlinearity **6**, 277 (1993), `chao-dyn/9303016`.
[27] M. Hénon, *Generating Families in the Restricted Three-Body Problem II. Quantitative Study of Bifurcations* (Springer, New York, 2001).
[28] J.-A. Sepulchre and R. Mackay, Nonlinearity **10**, 679 (1997).
[29] O. E. Rössler, Phys. Lett. A **57**, 397 (1976).
[30] M. Hénon and C. Heiles, Astron. J. **69**, 73 (1964).
[31] V. Szebehely, *Theory of Orbits* (Academic Press, New York, 1967).
[32] Y. Kuramoto and T. Tsuzuki, Progr. Theor. Physics **55**, 365 (1976).
[33] G. I. Sivashinsky, Acta Astr. **4**, 1177 (1977).